# Thinking Faster and Slower: An Agent's Cognitive Repertoire

Marieke van Otterdijk[1], Nathan Lloyd[2] and Peter R. Lewis[2]

[1]Vrije Universiteit Amsterdam, Amsterdam, The Netherlands
[2]Ontario Tech University, Oshawa, Canada
[1]m.t.h.van.otterdijk@vu.nl

**Abstract**

This work examines the limitations of artificial intelligence inspired by dual-process models—those that sharply divide cognition into fast, intuitive processes and slow, deliberative ones. In response, we propose an alternative framework that introduces a component for cognitive systems capable of representing a continuum of mental activity; the cognitive repertoire. Rather than relying on rigidly separated systems, we allow for faster and slower activity to operate in parallel on different time scales, providing fluid, context-sensitive behavior.

## Introduction

The field of artificial intelligence has long distinguished between fast, reactive systems and slow, deliberative ones, often contrasting heuristic-based responses with rule-based analytical reasoning (Bonnefon and Rahwan, 2020). This distinction parallels long-standing debates between symbolic and connectionist approaches, with the former emphasizing structured representations, logic, and explicit reasoning, and the latter favoring distributed, sub-symbolic computation, as seen in neural networks. Bridging the gap between these two approaches has been discussed as "*the* key" to artificial general intelligence (Goertzel, 2012). Since that claim, the field has advanced significantly—fueled by an explosion of available data, greater computational resources, and the growing influence of dual-process theories such as those outlined in *Thinking, Fast and Slow* (Kahneman, 2011)—all of which have introduced new tools and perspectives for designing psychologically inspired artificial intelligence. Indeed, several learning, processing, and decision-making models underpinned by dual-process theories have been proposed (Kahneman, 2011; Gigerenzer, 2024; Plessner et al., 2011). Kahneman and Tversky's work, in particular, is touted as an avenue toward stronger AI (Ganapini et al., 2022), leading to systems that are generalizable, adaptable, robust, explainable, and capable of causal analysis, abstraction, and common sense, among others.

Briefly, and while there are multiple dual-process theories (Evans, 2011; Plessner et al., 2011), the commonly understood dual-process theories suggest that decision-making involves a combination of an intuitive, experiential, or tacit system and an analytical, rational, or deliberate system. In the dual processing theory by Kahneman (2011); Plessner et al. (2011); Betsch (2008), the intuitive system is called System 1, and the logical system is called System 2. System 1 is automatic relying on emotions and impressions, highly prone to biases. In contrast, System 2 requires focus and attention, and is rational. Though often framed in terms of speed, autonomous vs. controlled is a stronger descriptor (Evans, 2011). Described by Lindström et al. (2022), this has come to mean for some, a pursuit of hybrid, or neuro-symbolic systems (Booch et al., 2021), or, as more recently demonstrated with large language 'reasoning' models, the use of mechanisms that route requests to specialized processes. In many cases, it is common to attribute 'fast' (automatic) to machine learning and deep learning (LeCun et al., 2015), and 'slow' (controlled) to logical reasoning. Indeed, if deep learning (LeCun et al., 2015) is one such basis for fast and intuitive responses, then its complement is symbolic reasoning with explicit representation, causal models, rules, and reasons (Bonnefon and Rahwan, 2020).

There are some fallacies associated with the dual-process theory (Evans, 2012) as well as some critiques (Gigerenzer, 2024). Gigerenzer contends that while opposing cognitive processes exist, they do not correspond to two distinct systems; instead, we may use the terms Type 1 and Type 2 processing (Evans, 2011). The dual-processing system models are generally critiqued for being too shallow in the use of fast and slow, holding false equivalence for 'fast' and neural, and failing to recognize the bi(multi)-directional influence of the reasoning systems (Lindström et al., 2022). While cognitively inspired algorithms may enable computers to perform tasks that minds can do (Boden, 2016, p.1), they can introduce restrictive assumptions. Furthermore, Da Silva (2023) highlights the criticism that dual-processing theories oversimplify cognition, as many mental processes involve the integration of both Type 1 and Type 2 processing. We expand upon this, avoiding the crisp binary mind and dual-system views, and, through our proposed architecture, posit a continuum between *fast* and *slow*; a cognitive repertoire.

## A Continuum of Cognitive Processes

We propose a cognitive architecture capable of completing reasoning processes that may be described as "faster" or "slower", or as requiring less or more information, across different time scales, with components for integration and reconciliation. This architecture receives input in the form of direct sense data, context maintained from the prior time step, alongside memory, drivers, and goals. An overview of the elements of the cognition layer is shown in Figure 1.

The layer consists of a *Cognitive Repertoire*, an adaptive toolbox for reasoning strategies (e.g., heuristics, learning models, or logic-based reasoning). A key function of the repertoire is to select processes suited to the incoming input; we allow for many processes to spin up concurrently. We emphasize that the processes which are activated are those well-suited to their environment, they are ecologically rational (Mata et al., 2012), which is not about 'rational' or 'irrational' in the dual-processing sense, rather the adaptive fit relative to the environment (Gigerenzer, 2004). These processes operate on a *Cognitive Process Spectrum*, a temporal window that allows multiple cognitive processes to run in parallel, enabling both faster and slower forms of reasoning to unfold. Examples of processes that may operate in parallel include heuristics or associations, which may be fully autonomous, estimation, which is slower with medium to high control, and task planning, which is slower still and demands high mental effort or control.

A *Satisfaction Regulator and Integrator* is necessary, given that multiple processes may operate in parallel, with varying degrees of interdependence and intent. This component of the system functions as a mental pause button and a coordination mechanism, capable of integrating interdependent processes for cohesive outputs, which has been inspired by the work of *merging* (Bellman and Krasne, 1983; Bellman, 1979), also referred to as "*blending*" (Stein et al., 1986). This stage enables sub-processes and dependencies to be resolved via direct aggregation and termination. Still, it may also allow the results of sub-processes to be fed back into the system as input, thereby avoiding unnecessary waiting cycles when "decisions" are unsatisfactory (incomplete, uncertain, unaligned, etc.) with the goals set by the agent. We align this component with the concept of '*satisficing*' (Simon, 1956). The Regulator waits for a satisfactory decision and passes that on to the output layer, where an action is put into motion.

Thus far, the proposed architecture and its continuum offer a more nuanced, flexible model: rather than the dual-processing approach. However, within this architecture, we also specify a reflective component inspired by recent work (Lewis and Sarkadi, 2024; Ganapini et al., 2022; Booch et al., 2021), and seminal work on meta-cognition and introspection (Sloman and Chrisley, 2003; Stanovich, 2009) in tripartite architectures, or a third Type of processing (Evans, 2009). While we visually distinguish reflective processes within Figure 1, reflection operates identically to the cognitive repertoire, containing processes that operate "faster" and "slower", but do so *only* at a meta-cognitive level. The slower, or controlled processes, may be akin to the executive control described by (Sloman and Chrisley, 2003). To avoid "stupid" reasoning (Sloman and Chrisley, 2003), the system may self-monitor and assess which type of reasoning suits the context, potentially terminating unsuitable processes and guiding the cognitive repertoire and integrator accordingly. The faster, or more autonomous processes, on the other hand, may be responsible for garbage collection, model retraining, and so on. These distinctions serve as exemplars, but reflective processes need not be discretely categorized, as our argument for a continuum may suggest. For example, both may be responsible for discovering cognitive shortcuts, causal rules, and abstractions, thereby adapting the repertoire or changing the evaluation criteria in the satisfaction regulator. This adaptation enables agents to exploit the new process, means of representation, or attend to salient features of the environment.

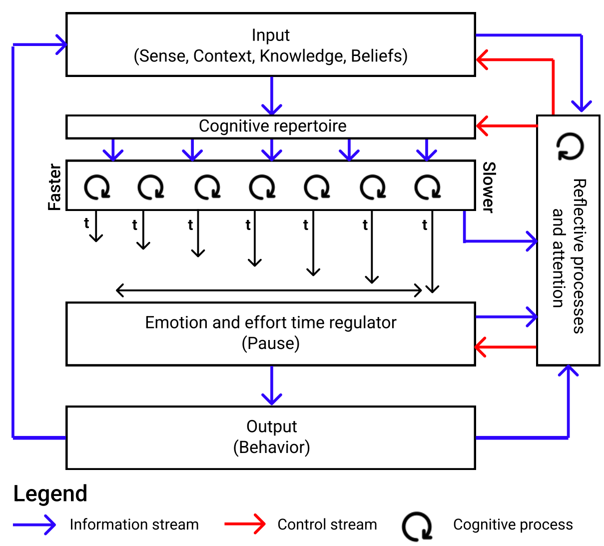


Figure 1: Overview of the cognitive repertoire architecture.

## Discussion

This paper proposes an architecture with a cognitive repertoire that allows for both faster and slower processes in agents. These processes are ecologically rational and capable of running in parallel. Their outputs are directed to a satisfaction regulator and integrator, which evaluates and integrates outputs. There are several critical factors for developers to consider, including the instantiation of the cognitive repertoire, goals, and knowledge, whether a reflective process can modify specific processes, and what sensory information is available. In future work, we will expand upon the architecture and opportunities for testing within virtual and physical agents.

## References


Bellman, K. (1979). *The Conflict Behavior of the Lizard, Sceloporus Occidentalis: And Its Implication for the Organization of Motor Behavior*. University of California, San Diego.

Bellman, K. and Krasne, F. B. (1983). Adaptive complexity of interactions between feeding and escape in crayfish. *Science*, 221(4612):779–781.

Betsch, T. (2008). The nature of intuition and its neglect in research on judgment and decision making. *Intuition in judgment and decision making*, pages 3–22.

Boden, M. A. (2016). *AI: Its nature and future*. Oxford University Press.

Bonnefon, J.-F. and Rahwan, I. (2020). Machine thinking, fast and slow. *Trends in Cognitive Sciences*, 24(12):1019–1027.

Booch, G., Fabiano, F., Horesh, L., Kate, K., Lenchner, J., Linck, N., Loreggia, A., Murgesan, K., Mattei, N., Rossi, F., et al. (2021). Thinking fast and slow in ai. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 15042–15046.

Da Silva, S. (2023). System 1 vs. system 2 thinking. *Psych*, 5(4):1057–1076.

Evans, J. (2012). Dual process theories of deductive reasoning: facts and fallacies. *The Oxford handbook of thinking and reasoning*, pages 115–133.

Evans, J. S. B. (2009). How many dual-process theories do we need? one, two, or many?

Evans, J. S. B. (2011). Dual-process theories of reasoning: Contemporary issues and developmental applications. *Developmental review*, 31(2-3):86–102.

Ganapini, M. B., Campbell, M., Fabiano, F., Horesh, L., Lenchner, J., Loreggia, A., Mattei, N., Rossi, F., Srivastava, B., and Venable, K. B. (2022). Thinking fast and slow in ai: The role of metacognition. In *International Conference on Machine Learning, Optimization, and Data Science*, pages 502–509. Springer.

Gigerenzer, G. (2004). Fast and frugal heuristics: The tools of bounded rationality. *Blackwell handbook of judgment and decision making*, 62:88.

Gigerenzer, G. (2024). *The intelligence of Intuition*. Cambridge University Press.

Goertzel, B. (2012). Perception processing for general intelligence: Bridging the symbolic/subsymbolic gap. In *International Conference on Artificial General Intelligence*, pages 79–88. Springer.

Kahneman, D. (2011). *Thinking fast and slow*. Penguin Books.

LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *nature*, 521(7553):436–444.

Lewis, P. R. and Sarkadi, Ş. (2024). Reflective artificial intelligence. *Minds and Machines*, 34(2):14.

Lindström, A. D., Mackay, W. E., and Dignum, V. (2022). Thinking fast and slow in human-centered ai. In *FSS 2022-AAAI Fall symposium Series-Thinking Fast and Slow and Other Cognitive Theories in AI*, pages 3–pages.

Mata, R., Pachur, T., Von Helversen, B., Hertwig, R., Rieskamp, J., and Schooler, L. (2012). Ecological rationality: a framework for understanding and aiding the aging decision maker. *Frontiers in neuroscience*, 6:19.

Plessner, H., Betsch, C., and Betsch, T. (2011). *Intuition in judgment and decision making*. Psychology Press.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological review*, 63(2):129.

Sloman, A. and Chrisley, R. (2003). Virtual machines and consciousness. *Journal of consciousness studies*, 10(4-5):133–172.

Stanovich, K. E. (2009). Distinguishing the reflective, algorithmic, and autonomous minds: Is it time for a tri-process theory. *In two minds: Dual processes and beyond*, pages 55–88.

Stein, P., Camp, A. W., Robertson, G. A., and Mortin, L. I. (1986). Blends of rostral and caudal scratch reflex motor patterns elicited by simultaneous stimulation of two sites in the spinal turtle. *Journal of Neuroscience*, 6(8):2259–2266.